ELSEVIER

Evolution and Human Behavior 32 (2011) 334–342

Evolution and Human Behavior

Original Article

An experimental comparison of human social learning strategies: payoff-biased social learning is adaptive but underused

Alex Mesoudi*

School of Biological and Chemical Sciences, Queen Mary University of London, London, UK

Initial receipt 5 July 2010; final revision received 9 December 2010

Abstract

Analytical models have identified a set of social learning strategies that are predicted to be adaptive relative to individual (asocial) learning. In the present study, human participants engaged in an ecologically valid artifact-design task with the opportunity to engage in a range of social learning strategies: payoff bias, conformity, averaging and random copying. The artifact (an arrowhead) was composed of multiple continuous and discrete attributes which jointly generated a complex multimodal adaptive landscape that likely reflects actual cultural fitness environments. Participants exhibited a mix of individual learning and payoff-biased social learning, with negligible frequencies of the other social learning strategies. This preference for payoff-biased social learning was evident from the initial trials, suggesting that participants came into the study with an intrinsic preference for this strategy. There was also a small but significant increase in the frequency of payoff-biased social learning over sessions, suggesting that strategy choice may itself be subject to learning. Frequency of payoff-biased social learning predicted both absolute and relative success in the task, especially in a multimodal (rather than unimodal) fitness environment. This effect was driven by a minority of hardcore social learners who copied the best group member on more than half of trials. These hardcore social learners were also above-average individual learners, suggesting a link between individual and social learning ability. The lower-than-expected frequency of social learning may reflect the existence of information producer–scrounger dynamics in human populations.

Keywords: Conformity; Cultural evolution; Payoff-biased social learning; Random copying; Social learning

1. Introduction

When faced with novel tasks, individuals can arrive at adaptive solutions either via *individual learning* (e.g., trial-and-error or reinforcement learning) or via *social learning* (i.e., copying the solution of another individual). Social learning is widespread in the animal kingdom (Galef & Laland, 2005) and especially pronounced in humans (Herrmann, Call, Hernandez-Lloreda, Hare & Tomasello, 2007; Horner & Whiten, 2005). Mathematical analyses have found that social learning is often more likely to lead to adaptive solutions than individual learning (Aoki, Wakano & Feldman, 2005; Boyd & Richerson, 1985, 1988; Kameda & Nakanishi, 2003; Rogers, 1988). This was demonstrated in the recent "social learning strategies tournament" (Rendell et al., 2010), in which scholars and members of the public submitted learning strategies that competed against one another in an unknown and relatively complex environment. Strategy success was significantly and positively predicted by the frequency with which that strategy copied the choice of another strategy, underlining the adaptiveness of social learning.

In the social learning strategies tournament, submitted strategies were forced to copy indiscriminately, i.e., copy the choice of a randomly selected fellow strategy. However, theoretical analyses have shown that rather than copying indiscriminately, it is often (even more) adaptive to follow a specific social learning strategy (Laland, 2004). Particularly, adaptive social learning strategies identified so far include copying the most common solution in the group, or "conformity" (Boyd & Richerson, 1985; Henrich & Boyd, 1998), and copying the solution of highly successful individuals in the group, or "payoff bias" (Boyd & Richerson, 1985; McElreath et al., 2008; Schlag, 1998).

---

* Corresponding author. Biological and Experimental Psychology Group, School of Biological and Chemical Sciences, Queen Mary University of London, Mile End Road, London E1 4NS, United Kingdom.

E-mail address: a.mesoudi@qmul.ac.uk.

1090-5138/$ – see front matter © 2011 Elsevier Inc. All rights reserved.
doi:10.1016/j.evolhumbehav.2010.12.001

The adaptiveness of different social learning strategies is seldom directly compared in such models. It is still unclear (i) which social learning strategy (or combination of strategies) is most adaptive in different situations, relative both to other potential social learning strategies and to individual learning, and (ii) which social learning strategy (or strategies) people actually follow in different situations, which, due to cognitive or information-processing constraints, may not necessarily be the most adaptive strategy identified by formal analyses.

To address Point (ii), the predictions of theoretical analyses have begun to be empirically tested using laboratory experiments (Caldwell & Millen, 2008; Efferson, Lalive, Richerson, McElreath & Lubell, 2008; Kameda & Nakanishi, 2002, 2003; McElreath et al., 2008; McElreath et al., 2005; Mesoudi, 2008; Mesoudi & Whiten, 2008; Toelch et al., 2008). Regarding Point (i), the only study to date to have directly and simultaneously compared more than one form of social learning was conducted by McElreath et al. (2008), who found that people tend to copy the behaviour of other group members who had the highest payoff (payoff bias), falling back on copying the most common behaviour in the group (conformity) when payoffs were tied. However, because that study was set up to address the issue of detecting social learning in non-human species in the wild, the task employed was rather simple (a single dichotomous choice of planting one of two crops in a farm) and is unlikely to be representative of many actual instances of human social learning.

The present study aimed to address the question of which social learning strategy (or strategies) people follow, and the adaptiveness of these social learning strategies, in a complex task more representative of actual human social learning and previously verified against real-world data (see Mesoudi, 2008; Mesoudi & O’Brien, 2008a). This task involves participants in groups designing “virtual arrowheads” via a computer program. Like actual projectile points, virtual arrowheads consist of multiple continuous (length, width, thickness) and discrete (shape, colour) attributes each of which affects overall arrowhead effectiveness according to prespecified fitness functions (see below). Participants can test the effectiveness of their design on a simulated “hunt” during which they receive a payoff determined by the fitness functions. Arrowheads can be improved (i.e., payoffs increased) over successive hunts either via trial-and-error individual learning or by copying the design of another group member. A previous study using this task (Mesoudi & O’Brien, 2008a) successfully recreated patterns of variation documented in the archaeological record of the Great Basin region of the southwestern United States by Bettinger and Eerkens (1999), supporting its ecological validity.

A key aspect of the task is the complexity of the fitness environment in which learning takes place. Each arrowhead attribute is associated with a separate fitness function, with these separate fitnesses summed to obtain the overall arrowhead payoff. These fitness functions can be either unimodal, in which case there exists a single optimal arrowhead design located at the single optimum of each attribute, or bimodal, in which case multiple locally optimal designs exist of varying maximum fitness. The latter can be conceptualised as a *multimodal adaptive landscape* with each point in the landscape representing a different arrowhead design (i.e., combination of attribute values) and the height of the landscape representing the payoff of that design (cf. Wright, 1932). Peaks of varying height in the landscape represent locally optimal designs of different maximum payoffs. Multimodal adaptive landscapes are likely to be more representative of actual technological evolution than simple dichotomous choices (Boyd & Richerson, 1992; Mesoudi & O’Brien, 2008a). Specifically pertaining to projectile points, systematic testing of replica points demonstrates functional tradeoffs between different demand characteristics (Cheshier & Kelly, 2006). For example, thin and narrow points maximise penetration of prey hides, whereas wide and thick points maximise wound size and bleeding. These can be seen as two adaptive peaks in the projectile point design landscape.

Previous experimental studies using the virtual arrowhead task have partially supported the predictions outlined above concerning the adaptiveness of social learning, showing that payoff-biased social learning is both willingly adopted by a majority of participants in preference to individual learning and is more adaptive than individual learning in multimodal adaptive landscapes (Mesoudi, 2008; Mesoudi & O’Brien, 2008a). These previous studies, however, did not allow participants to engage in alternative social learning strategies. This shortcoming is addressed in the present study. As well as payoff-biased social learning (copying the arrowhead design of the most successful group member) and individual learning (directly modifying attributes without copying anyone else) as before, participants could engage in three additional social learning strategies that have previously received theoretical attention. *Conformity* allowed participants to copy the most common arrowhead attributes across the group (Henrich & Boyd, 1998). *Averaging* allowed participants to copy the arithmetic mean of the arrowhead attributes across the entire group, as modelled by Boyd and Richerson (1985; p. 72) as “blending transmission”. Finally, *random copying* allowed participants to copy a randomly selected group member’s arrowhead design. Random copying is assumed in several of the aforementioned analytical studies, as well as in the social learning strategies tournament, as a “default” social learning strategy, although it has been suggested by Bentley, Hahn and Shennan (2004) to account for various real-world cultural dynamics such as the distribution of first names, dog breeds and pottery decorations.

A previous agent-based simulation (Mesoudi & O’Brien, 2008b) compared the adaptiveness of these four strategies (payoff-bias, conformity, averaging and random copying) in the virtual arrowhead task. Groups of agents, simulating human participants, employed one of the four strategies and were separately compared with individual learning-only control agents. These simulations showed that payoff bias outcompetes the other strategies only under the assumption of a multimodal adaptive landscape. Under this assumption, pure individual learners get stuck on locally optimal, but globally sub-optimal, designs. Payoff-biased social learners, by copying high-scoring group members, can jump from low-fitness designs to the higher-fitness design of that successful group member. Conformity allows individuals to converge on the most popular design, but in the absence of payoff-related information this design is no more likely to be the globally optimal design than one of the locally optimal designs. Random copying similarly fails to identify the globally optimal design, with copiers converging on a randomly selected locally optimal design. The averaging strategy performs particularly badly in multimodal adaptive landscapes given that the average of two peaks in the landscape is likely to be located in a low-fitness adaptive valley. In a unimodal adaptive landscape, in contrast, there was little difference between strategies, as pure individual learners could easily find the single optimal arrowhead design without getting stuck on sub-optimal peaks, and other strategies subsequently copied these individual learners.

The present study featured a similar set-up to the agent-based model of Mesoudi and O'Brien (2008b). Participants could either learn individually or copy the design(s) of a group of pure individual learners using one of the four aforementioned social learning strategies. This occurred either in a multimodal or in a unimodal adaptive landscape. It is predicted that participants will behave adaptively by employing the payoff-biased social learning strategy, especially in a multimodal adaptive landscape where this strategy is particularly adaptive. An additional, open question is whether participants do this spontaneously from the outset, or whether they gradually learn that this strategy is optimal. Simulations (including that of Mesoudi & O'Brien, 2008b) typically assume, for simplicity, that strategies are fixed, but it may be that people initially experiment with different strategies before fixing on the most appropriate strategy for a particular task.

2. Methods

2.1. Design

An initial group of 10 participants served as individual learning-only "demonstrators" from whom subsequent participants could copy arrowhead designs. Demonstrators engaged in three seasons of hunting each comprising 30 hunts during which they could only improve their arrowhead designs via individual (asocial, trial-and-error) learning. Five were randomly assigned to a unimodal and five to a multimodal fitness environment. Forty-eight different participants (henceforth "learners") then took part in the main experimental session. Like the demonstrators, learners engaged in three seasons of hunting each comprising 30 hunts and were randomly assigned to either a unimodal or a multimodal fitness environment. Additionally, learners could, during each hunt (except the first), (a) see their rank (first to sixth) relative to the scores on the equivalent hunt of the five demonstrators who experienced the same fitness environment as they did (unimodal or multimodal), and (b) optionally choose to copy the arrowhead design(s) of those five demonstrators according to one of four social learning strategies: payoff-bias, random copying, conformity or averaging. Note that the latter was optional, and learners could choose to ignore social information and learn individually. Choosing one of the social learning strategies automatically changed that learner's current arrowhead according to that strategy and they received a payoff associated with that copied design. All participants were paid a base rate of £5 for taking part. In order to increase motivation, an additional £1 was paid to learners for every season in which they were ranked first relative to the demonstrators at the final (30th) hunt, giving a maximum payment of £8. Previous studies (Mesoudi 2008; Mesoudi & O'Brien 2008a) found high intrinsic levels of motivation in this task, so this minimal incentive mechanism was deemed sufficient.

2.2. Materials/task

The aim of the virtual arrowhead experiment was to design a "virtual arrowhead" and test this arrowhead design in a "virtual hunting environment", all via a computer program (programmed in Borland C++ Builder 6.0; code available upon request). Screenshots of the task are shown in Fig. 1. Initially, and in all subsequent individual learning trials, participants enter values for five arrowhead attributes (Fig. 1A): length, width and thickness (continuous attributes each ranging from 1 to 100 arbitrary units), shape and colour (each taking one of four discrete values). Once valid values of each attribute have been entered, participants can press a HUNT button to receive feedback in calories, ranging from 1 to 1000. An initial practice session of 10 hunts with no opportunity to copy, and which did not contribute to the final payment, was performed to allow the participants to familiarize themselves with the task. The experiment proper was divided into three seasons each comprising 30 hunts (trials). The stated objective was to find the optimal values of each attribute that would give the highest possible score (1000 calories). This score was calculated using fitness functions (see Appendix A), although participants were told nothing about these functions. There was also random error in the feedback, with the score seen by participants randomly taken from a normal distribution with a mean given by the fitness of their arrowhead design and a standard deviation of 5 calories. Participants were informed about the presence of random error in feedback, but not its magnitude. During each season, the running total of all calories received during every hunt up to that point was displayed on screen, giving a personal cumulative score, and learners were informed how

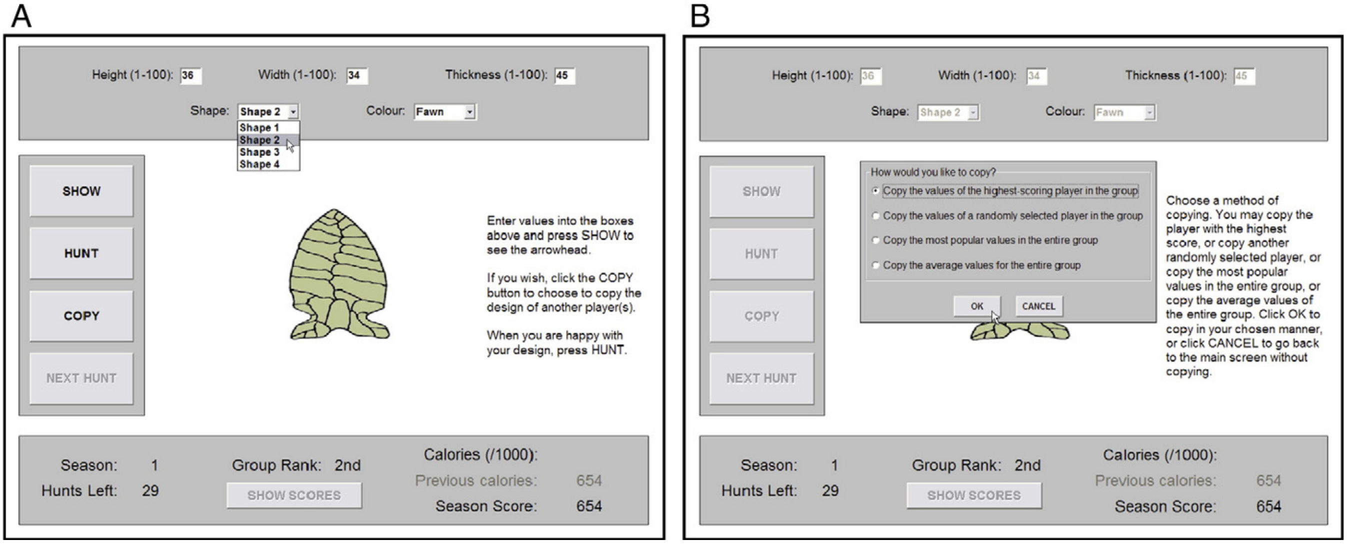

A

Height (1-100): 36
Width (1-100): 34
Thickness (1-100): 45
Shape: Shape 2
Colour: Fawn

Enter values into the boxes above and press SHOW to see the arrowhead.

If you wish, click the COPY button to choose to copy the design of another player(s).

When you are happy with your design, press HUNT.

Season: 1
Hunts Left: 29
Group Rank: 2nd
Calories (/1000):
Previous calories: 654
Season Score: 654

B

Height (1-100): 36
Width (1-100): 34
Thickness (1-100): 45
Shape: Shape 2
Colour: Fawn

How would you like to copy?
● Copy the values of the highest-scoring player in the group
○ Copy the values of a randomly selected player in the group
○ Copy the most popular values in the entire group
○ Copy the average values for the entire group

Choose a method of copying. You may copy the player with the highest score, or copy another randomly selected player, or copy the most popular values in the entire group, or copy the average values of the entire group. Click OK to copy in your chosen manner, or click CANCEL to go back to the main screen without copying.

Season: 1
Hunts Left: 29
Group Rank: 2nd
Calories (/1000):
Previous calories: 654
Season Score: 654

Fig. 1. Screenshots of the experimental task, where virtual arrowheads can be improved either (A) via individual trial and error, by directly modifying the values of height, width, thickness, shape and colour using the boxes along the top of the screen; or (B) by choosing one of four social learning strategies (payoff-bias, conformity, averaging or random copying).

they ranked relative to the equivalent demonstrators with respect to this cumulative score. At the start of each season the cumulative score was reset to 0 and the optimal values were changed, although optimal values did not change within each season. Participants were informed that optimal values changed between seasons and not during seasons.

The demonstrators played all three seasons with no opportunity to copy other participants, relying solely on individual trial-and-error learning (see Mesoudi & O'Brien 2008b for an analysis and simulation of the individual learning strategy employed by participants in this task). The learners could, during every hunt, optionally choose to copy the arrowhead design(s) of the demonstrators on the equivalent hunt according to one of four social learning strategies (Fig. 1B): (i) payoff-bias (described to participants as "Copy the values of the highest-scoring player in the group"), (ii) random copying ("Copy the values of a randomly selected player in the group"), (iii) conformity ("Copy the most popular values in the entire group") or (iv) averaging ("Copy the average values for the entire group"). If one of these was selected, then the existing arrowhead design was replaced with the chosen design as per the strategy, with no copying error. Learners had no foreknowledge of how effective a copied design would be before choosing a particular strategy (otherwise every strategy would effectively constitute a payoff-bias strategy). The payoff-bias strategy used players' cumulative season scores to determine the highest-scoring player, rather than scores on the immediately preceding hunt. The conformity and averaging strategies considered each attribute separately. For conformity, the continuous attributes (length, width, thickness) were divided into 10-unit intervals (1–10, 11–20, 21–30, etc.) and the mid-point of the most common interval across the five models was calculated. The most common values of the discrete attributes (shape and colour) were also calculated. Ties between two or more equally-most-common attributes were chosen from at random. For averaging, the arithmetic means of the continuous attributes across the group of five models were calculated. The discrete attributes were not modified by the averaging strategy because discrete values cannot be averaged. These methods of implementing the strategies were explained to the learners during the experiment via on-screen instructions.

2.3. Participants

Fifty-eight undergraduate students (43 female, mean age 20.43 years) from Queen Mary University of London participated for payment of between £5 and £8 (depending on performance, see above) and in some cases additional course credit. Ten of these served as demonstrators: five in the unimodal condition and five in the multimodal condition. The other 48 served as learners who had the opportunity to engage in the various forms of social learning: 24 in the unimodal condition and 24 in the multimodal condition.

2.4. Procedure

Because the design did not require entire groups to be present together, participants were run separately. Each participant sat at a computer and completed the experiment entirely via the on-screen interface. Participants were randomly allocated to unimodal or multimodal fitness environments. Participants were instructed not to talk or communicate with each other except via the computer program and not to write anything down. The entire study lasted 45–60 min. Informed consent was obtained from all participants and the study was approved by the Queen Mary University of London Research Ethics Committee.

2.5. Analysis

Nonparametric tests were conducted on the frequencies with which participants chose different strategies. Stepwise multiple regression analyses were performed with either score or rank as the dependent variable and strategy frequency (individual learning, payoff-bias, random copying, conformity and averaging) and fitness environment (unimodal or multimodal) as predictors.

3. Results

3.1. Learning strategy choices

Overall, the learners exhibited a mix of individual learning and payoff-biased social learning, with negligible amounts of the other social learning strategies. Table 1 shows the frequencies with which the 48 learners (i.e., excluding demonstrators) chose the different individual and social learning strategies over the three seasons of hunting. Overall, individual learning was the most common, chosen on 77.5% of hunts. Payoff-biased social learning was chosen on 18.8% of hunts, with negligible (approx. 1%) amounts of other strategies. Wilcoxon signed rank tests were used to compare the frequency of each strategy assuming a Bonferroni-corrected significance level of $p<.005$ (given 10 comparisons) and showed that individual learning was chosen significantly more often than payoff bias ($z=-4.92$, $p<.001$), random copying ($z=-6.03$, $p<.001$), conformity ($z=-5.97$, $p<.001$) and averaging ($z=-6.03$, $p<.001$). Similarly, payoff bias was chosen significantly more often than random copying ($z=-5.67$, $p<.001$), conformity ($z=-5.56$, $p<.001$) and averaging ($z=-5.70$, $p<.001$). Random copying, conformity and averaging did not significantly differ from one another.

Table 1 also shows an increase in the frequency of payoff-biased social learning from Season 1 (15.9% of hunts) to Season 2 (18.8%) to Season 3 (21.6%). Page's trend test (Page, 1963) for three repeated measures conditions showed this increase to be significant ($L=601$, $k=3$, $n=48$, $p<.01$). There was also a corresponding drop in the frequency of individual learning from Seasons 1 and 2 (78.5% and 78.9%) to Season 3 (75.2%), although this trend did not reach significance ($L=592$, $k=3$, $n=48$, $p>.05$). No significant trends were found for random copying ($L=572$, $k=3$, $n=48$, $p>.05$), conformity ($L=589$, $k=3$, $n=48$, $p>.05$) or averaging ($L=587$, $k=3$, $n=48$, $p>.05$).

Table 1
Percentage of hunts on which participants (excluding demonstrators) chose different learning strategies across three seasons of hunting (±S.E.)

Strategy	Individual learning	Payoff bias	Random copying	Conformity	Averaging
Season 1	78.5 (±3.5)	15.9 (±3.0)	0.7 (±0.3)	2.1 (±0.5)	2.9 (±1.3)
Season 2	78.9 (±4.2)	18.8 (±4.3)	0.7 (±0.5)	0.9 (±0.3)	0.7 (±0.3)
Season 3	75.2 (±4.5)	21.6 (±4.4)	1.4 (±0.8)	1.0 (±0.3)	0.7 (±0.4)
Total	77.5 (±3.8)	18.8 (±3.7)	1.0 (±0.4)	1.3 (±0.3)	1.4 (±0.6)

Individual learning and payoff bias significantly differed in frequency from each other strategy; random copying, conformity and averaging did not differ from one another.

3.2. Payoff differences between strategies

Multiple linear regression analysis was used to predict learners' success on the task, as measured by both cumulative score after all 30 hunts averaged across all three seasons, and final rank (1 to 6, with 1 indicating the top rank) relative to the demonstrators again averaged across all three seasons. Averaging these measures across seasons was justified by significant correlations between the three seasons in score (S1 vs. S2: $r=0.69$, $p<.001$; S1 vs. S3: $r=0.52$, $p<.001$; S2 vs. S3: $r=0.60$, $p<.001$) and rank (S1 vs. S2: $r_s=0.62$, $p<.001$; S1 vs. S3: $r_s=0.64$, $p<.001$; S2 vs. S3: $r_s=0.67$, $p<.001$). Predictors were the frequency of each learning strategy (individual learning, payoff-bias, random copying, conformity and averaging) as well as fitness environment (unimodal or multimodal). Variables were entered and removed in stepwise manner. For score, only frequency of payoff-biased social learning remained as a significant predictor (adj. $R^2=0.10$, $\beta=0.34$, $t=2.45$, $p<.02$) with all other predictors removed due to non-significance at $p<.05$. The relationship between payoff-bias frequency and score is shown in Fig. 2A. For rank, two predictors were significant (adj. $R^2=0.36$): frequency of payoff-bias ($\beta=-0.28$, $t=-2.39$, $p<.03$) and fitness environment ($\beta=-0.50$, $t=-4.19$, $p<.001$). Participants in the multimodal fitness environment had higher mean ranks than participants in the unimodal fitness environment. The relationship between payoff-bias frequency and rank is shown in Fig. 2B.

Inspection of Fig. 2 suggests that the significant relationship between payoff bias frequency and score/rank was driven by six high-scoring and high-ranking learners who had particularly high frequencies of payoff-biased social learning, each of whom chose to copy the best group member on more than half of the 29 social learning opportunities. The pre-Season 1 practice session, during which learners engaged in 10 hunts with no opportunity for social learning, gives an opportunity to see whether these six hardcore social learners were also better at individual learning, as suggested by Reader and Laland's (2002) finding that innovation and social learning are linked across species. Although limited by the small and unequal sample size, these six hardcore social learners did indeed have significantly higher cumulative scores during the 10 practice hunts (mean=7138) than the other 42 learners (mean=5684; $t_{(46)}=2.15$, $p<.04$).

3.3. Effect of fitness environment

The regression analyses above show that rank, but not score, was significantly predicted by fitness environment: learners in a multimodal environment had higher ranks than learners in a unimodal environment. This is because rank

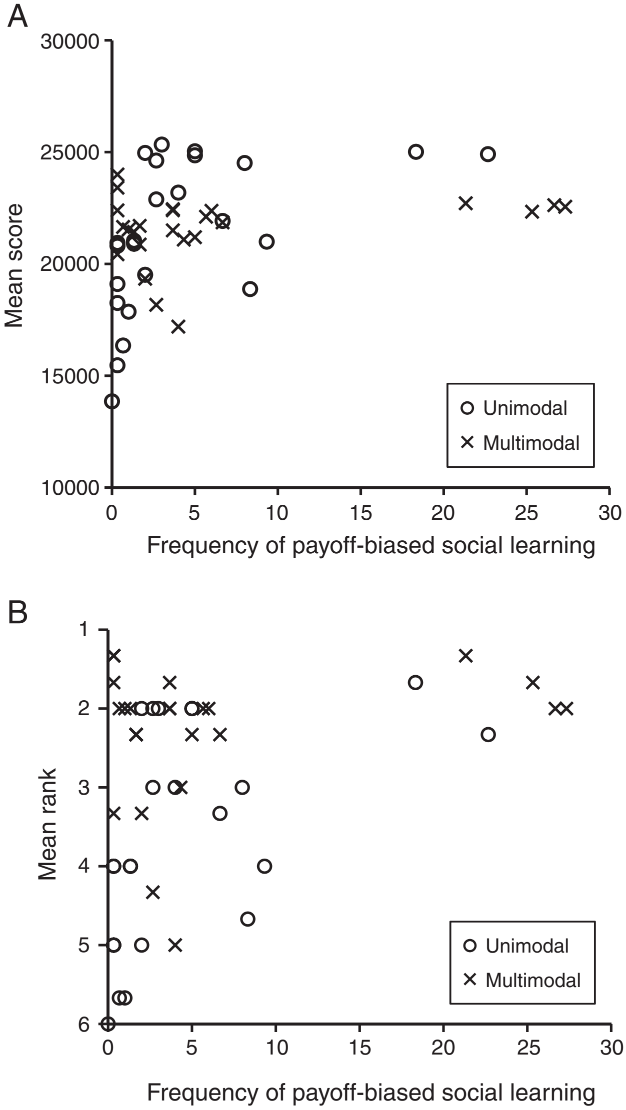

Fig. 2. Relationship between the number of hunts on which payoff-biased social learning was chosen (out of 29, averaged across all three seasons) and (A) mean final score and (B) mean final rank. Circles indicate unimodal participants; crosses indicate multimodal participants. Fitness environment (unimodal or multimodal) was a significant predictor of rank only.

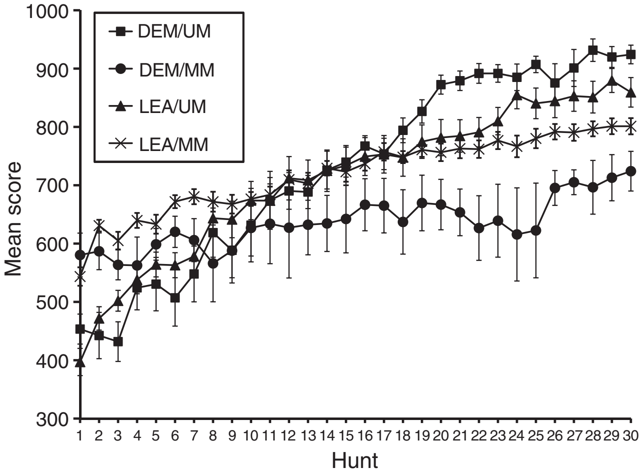

Fig. 3. Mean score over successive hunts, averaged across all three seasons. DEM=Demonstrators (who could not engage in social learning), LEA=learners (who could choose to engage in social learning), UM=unimodal fitness environment, MM=multimodal fitness environment. Error bars show standard error.

(but not score) was calculated relative to the group of asocial demonstrators. Multimodal demonstrators got stuck on locally optimal but globally suboptimal peaks in the multimodal landscape, thereby reducing their mean score. Multimodal learners, on the other hand, could escape their sub-optimal peaks by copying the designs of the highest-scoring group member who had located a higher-fitness peak. Unimodal demonstrators could easily locate the single peak in the landscape, allowing them to perform just as well as the unimodal learners. This scenario is supported by a direct comparison of scores for the different groups (Fig. 3). Participants (demonstrators and learners combined) in the unimodal environment started off at Hunt 1 with a significantly lower score than participants in the multimodal environment (UM mean=407, MM mean=550; $t_{(56)}$=5.78, p<.001) because the latter were more likely purely by chance to find themselves on one of the locally optimal fitness peaks, whereas unimodal participants were less likely to find themselves on the single peak in their fitness environment. By the final hunt, however, this was reversed and unimodal participants had a significantly higher mean score (UM=871, MM=788; $t_{(56)}$=3.23, p<.002). This is because multimodal participants get stuck on globally suboptimal fitness peaks, whereas unimodal participants all eventually converge on the single peak in their fitness landscape. The opportunity for social learning reduced this handicap of multimodal environments, as indicated by significantly higher final scores for multimodal learners (mean=801) than multimodal demonstrators (mean=724; $t_{(27)}$=2.26, p<.04). This latter difference was only marginally significant, possibly because not all of the learners learned socially (see Fig. 2). In contrast, there was no significant advantage of social learning for unimodal groups as indicated by the final score (UM demonstrators=924, UM learners=859; $t_{(27)}$= 1.15, p<.26). Finally, although there was a higher frequency of payoff-biased social learning in multimodal learners (mean=22.4%) than in unimodal learners (mean=15.2%), this difference was not significant according to a Mann–Whitney test (z=−0.69, n_1=n_2=24, p<.49).

4. Discussion

The aim of this study was to assess people's choices of learning strategy while engaged in a relatively complex and ecologically valid learning task. This task involved designing an arrowhead composed of multiple continuous and discrete attributes, with these attributes forming either a unimodal (single optimal arrowhead design) or a multimodal (multiple locally optimal arrowhead designs of varying maximum payoffs) fitness environment. An agent-based simulation analysis of this task (Mesoudi & O'Brien, 2008b) predicted that participants should engage in payoff-biased social learning (copying the design of the most successful group member) in preference to the other strategies (conformity, averaging, random copying, plus individual learning) especially in the multimodal fitness environment, as payoff bias allows learners to avoid getting stuck on sub-optimal peaks.

These predictions were partially supported. Participants exhibited a mix of individual learning and payoff-biased social learning, with negligible amounts of the other social learning strategies (conformity, averaging and random copying). The frequency of payoff bias was significantly associated with both absolute (score) and relative (rank) measures of success in the task. However, this effect was driven mainly by a small minority (12.5%) of learners who engaged in very high levels of payoff-biased social learning, with the rest of the learners engaging predominantly in individual learning. The shape of the adaptive landscape determined the success of payoff-biased social learning relative to asocial-learning-only demonstrators (as measured by rank): payoff-biased social learners only outperformed asocial demonstrators in the multimodal fitness environment because payoff bias allowed learners to escape low-fitness peaks; in unimodal fitness environments there was no advantage to social learning, because asocial demonstrators could easily find the single adaptive peak. However, this advantage did not influence the frequency of social learning: learners in multimodal environments were no more likely to engage in social learning (relative to individual learning) than learners in unimodal environments.

The finding that only a minority of learners consistently engaged in payoff-biased social learning (or social learning of any kind) replicates previous experimental findings related to conformity (Efferson et al., 2008; McElreath et al., 2005). Interestingly, the hardcore social learners in the present study were also better individual learners as assessed in a preliminary practice session, providing tentative within-species support for Reader and Laland's (2002) finding that social and individual learning abilities are linked across species. So rather than social learners falling back on a copying strategy in response to their poor individual learning ability, it seems that social learners were better learners in general. Whether this is, in turn, related to general cognitive abilities such as *g* remains a question for future research. Nevertheless, the finding that the majority of participants fail to learn socially, despite the adaptiveness of social learning in the task employed here, requires explanation. It may be that human populations constitute a mixed equilibrium comprising both information producers (i.e., individual learners) and information scroungers (i.e., social learners), as predicted by theoretical analyses (Barnard & Sibly, 1981; Giraldeau, Valone & Templeton, 2002; Rogers, 1988) and found in some experimental tasks (Kameda & Nakanishi, 2002, 2003), in which the lower costs of social learning are balanced by the need to track non-stationary environments via individual learning. Although individual learning was suboptimal in the experimental task employed here, participants may nevertheless be following an individual learning strategy because it is adaptive in their everyday lives. Further experiments employing non-stationary environments (e.g., peaks that move around in the adaptive landscape) and field studies of people's everyday learning behavior could address these issues. Broader cross-cultural differences may also be relevant here. Westerners, and particularly Western college students, tend to be highly individualistic relative to non-Westerners and non-college students (Henrich, Heine & Norenzayan, 2010), possibly explaining our Western college student participants' over-reliance on individual learning. Replications in non-Western populations may reveal higher levels of social learning.

Payoff-biased social learning was by far the most frequent social learning strategy from the very first hunt, dwarfing the frequencies of conformity, averaging and random copying. This indicates that payoff bias is intuitively appealing as a social learning strategy, replicating the findings of McElreath et al. (2008) and the predictions of previous simulations that payoff bias is uniquely adaptive in this specific task (Mesoudi & O'Brien, 2008b). This preference for payoff bias over other social learning strategies might be a product of genetic evolution and/or acquired during development, which future developmental studies with children might address. There was a significant trend for the frequency of payoff bias to increase over the three seasons, possibly indicating the flexible learning of strategy effectiveness over time. Perhaps extending the number of seasons/hunts would lead to increasingly higher frequencies, even in those participants who engaged predominantly in individual learning.

A concern with experimental tests of social learning, as for all experiments, is their lack of external validity. While previous experimental studies (e.g. Kameda & Nakanishi, 2002, 2003; McElreath et al., 2005) have typically used a rather artificial and highly simplified task with a single dichotomous choice (e.g., plant one of two kinds of crops or locate a rabbit in one of two places), the present study used a more complex and realistic task in which participants designed an artifact composed of multiple continuous and discrete attributes, with artifact fitness determined by a complex adaptive landscape. These artifact attributes (width, thickness, etc.) are the same as those modified by real-life hunter–gatherer arrowhead manufacturers and measured by archaeologists (e.g., Bettinger & Eerkens, 1999). Yet this task still contained several strong simplifying assumptions, particularly in the implementation of the social learning strategies. In real life, the identity of the most successful demonstrator may be ambiguous and determined indirectly through markers of prestige (Henrich & Gil White, 2001), rather than directly provided to learners as in the present study's implementation of payoff bias. The averaging and conformity strategies also provided learners with direct and accurate information about the average and most common arrowhead design, respectively, information which in real life would have to be calculated by the learner from often ambiguous environmental input. In a sense, these simplifications make the underuse of social learning in the present study all the more surprising, given that the implementation here removed many of the ambiguities and inaccuracies of such strategies and should have increased their use. On the other hand, perhaps these strategies were consequently so unfamiliar to the participants compared to how they are implemented in real life that the participants misapplied them. While simplified experimental tasks such as the one employed here are useful in allowing researchers to test clear and unambiguous theoretical predictions, future models and experiments might increase the complexity of learning tasks further to address these higher-order issues and further explore the apparent mismatch between theory and data.

The multimodal adaptive landscape simulated in the present experiment is likely to be representative of much human cultural evolution, particularly technological evolution (Boyd & Richerson, 1992; Mesoudi & O'Brien, 2008a). Many problems are likely to have multiple solutions of varying payoffs, with constant uncertainty over whether the current solution is globally or only locally optimal. Recent studies suggest that chimpanzees find it difficult to switch from adequate solutions to better solutions, even when those better solutions are clearly modelled (Hrubesch, Preuschoft & van Schaik, 2009; Marshall-Pescini & Whiten, 2008). This "switching" has been suggested to underlie humans' unique capacity for cumulative cultural evolution, in which successive improvements are built up over time (Whiten, McGuigan, Marshall-Pescini & Hopper, 2009). Payoff-biased social learning, which appears to be spontaneously employed by a significant minority of participants in the present study, may be a key mechanism driving this cumulative cultural evolution.

Acknowledgments

I am grateful to Jeremy Kendal, Kevin Laland, three anonymous reviewers and editor Robert Kurzban for useful comments on earlier versions of the paper.

Appendix A. Fitness equations

The fitness of an arrowhead, W ($0 \leq W \leq 1000$), was determined by the participant's entered attribute values of length, width, thickness and shape (a fifth attribute, colour, had no effect on fitness), according to Eq. (1):

$$W = 1000(0.275\ W_{\mathrm{L}} + 0.25\ W_{\mathrm{W}} + 0.35\ W_{\mathrm{T}} + 0.125\ W_{\mathrm{S}}) \quad (1)$$

where W_{L}, W_{W}, W_{T} and W_{S} are the individual fitness contributions of length, width, thickness and shape, respectively, each ranging from 0 to 1. Hence, these attributes are differentially weighted in descending order of importance of thickness, length, width and shape. Participants received feedback based only on the single number, W, and were not informed of separate fitness contributions of the attributes. For the multimodal environment, the three continuous attributes, length, width and thickness, each had bimodal fitness functions comprising two overlapping normal distributions with means O_1 and O_2, representing the global optimum and local optimum, respectively. The global optimum gave the maximum fitness contribution of the attribute (e.g., W_{L}=1), while the local optimum gave two-thirds that value (e.g., W_{L}=0.66). Each attribute had different optima, e.g., for length, O_{L1} and O_{L2}, for width O_{W1} and O_{W2}, and so on. The fitness contribution from each of these optima (using length as an example), W_{L1} and W_{L2}, was given by Eqs. (2) and (3):

$$W_{\mathrm{L1}}(X_{\mathrm{L}}, O_{\mathrm{L1}}) = P_1 exp\left[-(X_{\mathrm{L}}/100 - O_{\mathrm{L1}}/100)^2/(2\sigma)\right] \quad (2)$$

$$W_{\mathrm{L2}}(X_{\mathrm{L}}, O_{\mathrm{L2}}) = P_2 exp\left[-(X_{\mathrm{L}}/100 - O_{\mathrm{L2}}/100)^2/(2\sigma)\right] \quad (3)$$

where X_{L} is the participant's entered value of length, P_1 and P_2 are the maximum fitness contributions of the two optima (P_1=1, P_2=0.66), and σ is the standard deviation of the normal distribution (σ=0.025). The overall fitness contribution of length, W_{L}, is then the greater of the two values W_{L1} and W_{L2} according to Eq. (4):

$$W_{\mathrm{L}} = \begin{cases} W_{\mathrm{L1}} & \text{if } W_{\mathrm{L1}} > W_{\mathrm{L2}} \\ W_{\mathrm{L2}} & \text{if } W_{\mathrm{L2}} > W_{\mathrm{L1}} \end{cases} \quad (4)$$

Equivalent expressions for the fitness contributions of width and thickness are obtained by substituting different subscripts. For the participants in the unimodal environment, fitness was given by a single normal distribution [i.e., Eqs. (1) and (2) only, ignoring Eqs. (3) and (4)]. The fitness contribution of shape was identical for both unimodal and multimodal participants and was determined by ranking each of the four discrete shapes (Shape 1, Shape 2, Shape 3 and Shape 4) into a predetermined random order, with the highest ranked value giving a fitness contribution of W_{S}=1, the second W_{S}=0.9, the third W_{S}=0.66 and the last W_{S}=0.33. Finally, the actual feedback shown to participants, W', was calculated by adding random error to W. Hence, W' was drawn randomly from a normal distribution with mean W and standard deviation ε, where ε=5.

References

Aoki, K., Wakano, J. Y., & Feldman, M. W. (2005). The emergence of social learning in a temporally changing environment: A theoretical model. *Current Anthropology*, *46*, 334–340.

Barnard, C. J., & Sibly, R. M. (1981). Producers and scroungers: A general model and its application to captive flocks of house sparrows. *Animal Behaviour*, *29*, 543–550.

Bentley, R. A., Hahn, M. W., & Shennan, S. J. (2004). Random drift and culture change. *Proceedings of the Royal Society B*, *271*, 1443–1450.

Bettinger, R. L., & Eerkens, J. (1999). Point typologies, cultural transmission, and the spread of bow-and-arrow technology in the prehistoric Great Basin. *American Antiquity*, *64*, 231–242.

Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. Chicago, IL: University of Chicago Press.

Boyd, R., & Richerson, P. J. (1988). An evolutionary model of social learning: The effects of spatial and temporal variation. In T. Zentall, & B. G. Galef (Eds.), *Social Learning* (pp. 29–49). Hillsdale, NJ: Erlbaum.

Boyd, R., & Richerson, P. J. (1992). How microevolutionary processes give rise to history. In M. Nitecki, & D. V. Nitacki (Eds.), *History and evolution* (pp. 179–209). Albany: State University of New York Press.

Caldwell, C., & Millen, A. E. (2008). Experimental models for testing hypotheses about cumulative cultural evolution. *Evolution and Human Behavior*, *29*, 165–171.

Cheshier, J., & Kelly, R. L. (2006). Projectile point shape and durability: The effect of thickness:length. *American Antiquity*, *71*, 353–363.

Efferson, C., Lalive, R., Richerson, P. J., McElreath, R., & Lubell, M. (2008). Conformists and mavericks: The empirics of frequency-dependent cultural transmission. *Evolution and Human Behavior*, *29*, 56–64.

Galef, B. G., & Laland, K. N. (2005). Social learning in animals: Empirical studies and theoretical models. *BioScience*, *55*, 489–499.

Giraldeau, L. A., Valone, T. J., & Templeton, J. J. (2002). Potential disadvantages of using socially acquired information. *Philosophical Transactions of the Royal Society of London B*, *357*, 1559.

Henrich, J., & Boyd, R. (1998). The evolution of conformist transmission and the emergence of between-group differences. *Evolution and Human Behavior*, *19*, 215–241.

Henrich, J., & Gil White, F. J. (2001). The evolution of prestige. *Evolution and Human Behavior*, *22*, 165–196.

Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, *33*, 61–135.

Herrmann, E., Call, J., Hernandez-Lloreda, M. V., Hare, B., & Tomasello, M. (2007). Humans have evolved specialized skills of social cognition: The cultural intelligence hypothesis. *Science*, *317*, 1360.

Horner, V., & Whiten, A. (2005). Causal knowledge and imitation/emulation switching in chimpanzees (*Pan troglodytes*) and children (*Homo sapiens*). *Animal Cognition*, *8*, 164–181.

Hrubesch, C., Preuschoft, S., & van Schaik, C. (2009). Skill mastery inhibits adoption of observed alternative solutions among chimpanzees (*Pan troglodytes*). *Animal Cognition*, *12*, 209–216.

Kameda, T., & Nakanishi, D. (2002). Cost–benefit analysis of social/cultural learning in a nonstationary uncertain environment. *Evolution and Human Behavior*, *23*, 373–393.

Kameda, T., & Nakanishi, D. (2003). Does social/cultural learning increase human adaptability? Rogers' question revisited. *Evolution and Human Behavior*, *24*, 242–260.

Laland, K. N. (2004). Social learning strategies. *Learning and Behavior*, *32*, 4–14.

Marshall-Pescini, S., & Whiten, A. (2008). Chimpanzees (*Pan troglodytes*) and the question of cumulative culture: An experimental approach. *Animal Cognition*, *11*, 449–456.

McElreath, R., Bell, A. V., Efferson, C., Lubell, M., Richerson, P. J., & Waring, T. M. (2008). Beyond existence and aiming outside the laboratory: Estimating frequency-dependent and pay-off-biased social learning strategies. *Philosophical Transactions of the Royal Society B*, *363*, 3515–3528.

McElreath, R., Lubell, M., Richerson, P. J., Waring, T. M., Baum, W., Edsten, E., et al. (2005). Applying evolutionary models to the laboratory study of social learning. *Evolution and Human Behavior*, *26*, 483–508.

Mesoudi, A. (2008). An experimental simulation of the 'copy-successful-individuals' cultural learning strategy: Adaptive landscapes, producer–scrounger dynamics and informational access costs. *Evolution and Human Behavior*, *29*, 350–363.

Mesoudi, A., & O'Brien, M. J. (2008a). The cultural transmission of Great Basin projectile point technology: I. An experimental simulation. *American Antiquity*, *73*, 3–28.

Mesoudi, A., & O'Brien, M. J. (2008b). The cultural transmission of Great Basin projectile point technology: II. An agent-based computer simulation. *American Antiquity*, *73*, 627–644.

Mesoudi, A., & Whiten, A. (2008). The multiple roles of cultural transmission experiments in understanding human cultural evolution. *Philosophical Transactions of the Royal Society B*, *363*, 3489–3501.

Page, E. B. (1963). Ordered hypotheses for multiple treatments: A significance test for linear ranks. *Journal of the American Statistical Association*, *58*, 216–230.

Reader, S. M., & Laland, K. N. (2002). Social intelligence, innovation, and enhanced brain size in primates. *Proceedings of the National Academy of Sciences*, *99*, 4436–4441.

Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M., et al. (2010). Why copy others? Insights from the social learning strategies tournament. *Science*, *328*, 208–213.

Rogers, A. R. (1988). Does biology constrain culture? *American Anthropologist*, *90*, 819–831.

Schlag, K. H. (1998). Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *Journal of Economic Theory*, *78*, 130–156.

Toelch, U., van Delft, M. J., Bruce, M. J., Donders, R., Meeus, M. T. H., & Reader, S. M. (2008). Decreased environmental variability induces a bias for social information use in humans. *Evolution and Human Behavior*, *30*, 32–40.

Whiten, A., McGuigan, N., Marshall-Pescini, S., & Hopper, L. M. (2009). Emulation, imitation, over-imitation and the scope of culture for child and chimpanzee. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*, 2417–2428.

Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding and selection in evolution. *Proceedings of the Sixth International Congress of Genetics*, *1*, 356–366.